A model of long term learning: Integration of knowledge and knowledge compilation

Kurt VanLehn
Learning Research and Development Center
University of Pittsburgh
Pittsburgh, PA   15260

19950511 070

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 5-5-95	3. REPORT TYPE AND DATES COVERED Final Technical Report 7/1/92-11/15/93

4. TITLE AND SUBTITLE A model of long term learning: Integration of knowledge acquisition and knowledge compilation	5. FUNDING NUMBERS N00014-92-J-1945
6. AUTHOR(S) Kurt VanLehn	

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Kurt VanLehn 819 Learning Research and Development Center 3939 O'Hara St. University of Pittsburgh, Pittsburgh, PA 15260	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of the Navy Office of the Chief of Naval Research 800 N. Quincy St., Code 1511:LG Arlington, VA 22217-5000	10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

N/A

12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution unlimited; approved for public release.	12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

This report describes progress toward constructing a unified, computer-based model of all the major phenomena in cognitive skill acquisition. An extensive review of the literature was completed and published, along with several new analyses of important but hitherto neglected aspects of cogniitve skill. However, the computational embodiment of the theory was only partially implemented during the 15 months of funding, and more work is needed.

14. SUBJECT TERMS cognitive skill acquisition, knowledge compilation, computer-based model of cognitive skill acquisition	15. NUMBER OF PAGES 7
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT unclassified	20. LIMITATION OF ABSTRACT unclassified

Final Report:
A model of long term learning: Integration of knowledge acquisition and knowledge compilation

Kurt VanLehn

Computer Science Department
University of Pittsburgh
Pittsburgh, PA 15260
VanLehn@cs.pitt.edu
(412) 624-7458 fax (412)624-9149

May 5, 1995

Abstract

This report describes progress toward constructing a unified, computer-based model of all the major phenomena in cognitive skill acquisition. An extensive review of the literature was completed and published, along with several new analyses of important but hitherto neglected aspects of cognitive skill. These begin to articulate a new and exceedingly simple theory of cognitive skill. However, the computational embodiment of the theory was only partially implemented during the 15 months of funding, and more work is needed.

1 Objectives

It is useful and traditional to view cognitive skill acquisition as having three phases. The *early* phase consists of studying expository material, such as a text or a lecture, that teaches the basic principles and procedures of the skill. The *intermediate* phase initiates learning how to put these basic ideas into practice by solving problems with them. The main difference between the early and intermediate phases is that during the intermediate phase, the student focusses mainly on solving a problem or explaining an example with ocassional interruptions to refer to the text, ask a question or reflect on the task, whereas during the early phase, the student focusses mainly on understanding the text or teacher, and is not actively involved in trying to solve a problem or study an example. The *late* phase begins when the student can turn out error-free solutions on a regular basis, thus indicating that they have mastered the conceptual material of the task domain. They may still make unintentional errors (slips). The frequency of slips and the time to to a task decrease slowly with practice during late phase.

The objective of this grant was to extend CASCADE to cover late-phase effects. CASCADE was developed under an earlier ONR grant as a model of the intermediate phase of cognitive skill acquisition, and of self-explanation effect in particular (VanLehn et al., 1992; VanLehn and Jones, 1993a). The basic idea was to add a simple model of memory and demonstrate

that it could handle all the major phenomena of the intermediate and late phases. These phenomena would be demonstrated by modeling the acquisition of physics expertise.

The project was divided into four subtasks. Each subtask will be discussed in turn.

2 Developing an architecture

The original CASCADE system was built in Prolog with no explicit model of memory. It had grown quite baroque as it evolved, and we doubted that we could add a model of memory in any simple way. Thus, we searched for an architecture (i.e., a programming language with an integrated model of memory) that would make modeling physics easy.

We considered Actr, Soar, Ops and several other off-the-shelf architectures. None were adequate, basically because they used an attribute-value representation and matching, whereas physics is most simply represented with a clausal representation and unification. Thus, we developed a production system used Prolog-like clauses as its working memory elements and unification instead of matching in the production system interpreter. With John Anderson's blessing, we copied the production strengthening code from Actr.

In the course of this development, we spent considerable time evaluating the evidence for the procedural-declarative distinction. Leaving neurological evidence aside, there are basically two main pieces of evidence for the distinction. One is that the learning curves for individual production rules fit a power law beautifully, except for the first trial, which takes considerably longer than it should (Anderson et al., 1989; Bovair et al., 1990). The extra latency has been taken as evidence of a process that converts declarative to procedural knowledge. However, more careful experimental work indicates that this extra latency isn't really there (Anderson and Fincham, 1994).

The second type of evidence is the use-specificity of transfer. Practice up to a certain point (several dozen trials) causes a reduction in time-to-mastery of the transfer task, but practice after than point causes no further time savings. The initial practice is taken as strengthening declarative knowledge, which is shared between the tasks, and thus the initial practice benefits the transfer task. The later practice is taken as strengthening only the procedural knowledge, which is not shared between the tasks, and thus the later practice does not benefit the transfer task (Singley and Anderson, 1989). Examining protocols of learning during the first dozen or so uses of a rule (VanLehn, 1995d) convinced us that there is a conversion of knowledge going on, but it is conversion of some vague, poorly understood general instructions into a complete, fully debugged, operational procedure for solving problems. This is an entirely conscious process. One can hear the learning events in protocols as subjects debug their skill. The initial debugging transfers to new tasks because students are primarily concerned with understanding parts of the instructional material that the they didn't understand the first time they read it. The later practice does not transfer because it mostly results in speed-up of the task-specific procedure that the subjects have constructed, and removal of a few bugs that are specific to that procedure. Thus, we think that there is a shift from a superficial understanding of the instruction to an operational understanding, but both forms of knowledge are consciously accessible and the conversion is a deliberate process. In contrast, the Actr position is that the operational form of knowledge is not consciously accessible and the conversion is automatic.

We also spent considerable time pondering a second distinction, which is the difference between rule-based and schema-based (or case-based) architectures. The difference is primarily one of the size of the unit of knowledge retrieved. Experts do seem to retrieve large

pieces of knowledge all at once, but novice retrieve knowledge in small pieces. The schema-based architectures can model experts, but they have a difficulty time modeling novices. The rule-based architectures are good for modeling novices. Given some mechanism of association (e.g., spreading activation), they should be able to handle schema effects as well. Although this has not been demonstrated computationally, it seemed worth trying.

In short, re-examining the literature on cognitive architectures convinced us that a good old-fashioned production system would work best as an architecture for the new CASCADE. However, the architecture would be non-standard in two respects. First, it would have some kind of strengthening and context effect parameters associated with each rule. These would get updated with each application of the rule, and would affect its probability of retrieval. Second, working memory elements and rules were treated exactly the same way. They have the same kind of memory parameters. Learning a new rule is done by creating a working memory element that has the syntax of a rule.

Although the architecture sounds simple, its implementation turned out to be overly complex. The review of the architecture literature went on simultaneously with the development of the implementation. As a consequence, the implementation embeds some assumptions that were once thought to be correct, but are now no longer believed. However, it works, so we have not dared (yet) to throw it out and build a simpler one that conforms to present beliefs.

3   Review of the cognitive skill acquisition literature

Given the goal of accounting for all the major findings in cognitive skill acquisition, it was important to update the list of findings from the one presented in the grant proposal, which was based on VanLehn (1990). The resulting review, which is in press (VanLehn, 1995b), concluded that there are four basic sets of findings:

- *Practice/Transfer.* This group includes research on the power law of practice, automatization and the identical elements theory of transfer. The most recent relevant work has tried to understand the use-specificity of practice: why does increasing the practice on a training task decreases the amount of transfer to a transfer task (Anderson and Fincham, 1994)? Other, less relevant work has tried to settle basic questions of memory architecture using practice effects. For instance, does retrieval get easier because memory items get stronger, or there are more duplicate copies of them (Logan, 1988)?

- *Expert/Novice.* No major new phenomena have been discovered since the 1990 review, but Ericsson has proposed a good unifying framework based on the idea that practice causes experts to adapt their knowledge to the task domain (Ericsson and Lehmann, 1996). In particular, for task domain such as physics where one has all the information needed to solve a problem at the start, experts develop the ability to plan solutions mentally. This accounts for their ability to classify problems on the basis of the problems' solutions rather than their surface features (Chi et al., 1981), their changes in strategy (or lack thereof) (Larkin, 1983; Priest and Lindsay, 1992), and their improved memory of intermediate states (Chase and Simon, 1973).

- *Good/Poor Learners.* This group of studies includes investigations of the self-explanation effect, reflection and tutoring techniques such as feedback and mastery-based advance-

ment. This is mostly new work. The notion of a learning event seems crucial for explaining most of the phenomena (VanLehn, 1995b).

- *Schema acquisition.* This group of studies, most of which are fairly recent, investigates the way students use examples to help them solve problems, and thereby develop general solution techniques (schemas). This literature is a confusing mess until one notices that learning is being measured in two ways. In the majority of the studies, the examples and problems are isomorphic, so copying a solution from the example to the problem will get the problem right. Success thus depends mostly on retrieval and generalization. The other studies use problems that are not isomorphic to the examples. They generally find very little learning, and only under special circumstances which tend to promote self-explanation and other effective studying strategies (VanLehn, 1995b).

As it turns out, the key to understanding the more recent literature was conducting a fine-grained protocol analysis of the process of referring to examples while studying problems (VanLehn, 1995a). This process, which is often called analogical problem solving (somewhat misleadingly, I might add), is involved in the schema acquisition studies, but those studies used only outcome measures instead of protocol analysis or other process measures.

The protocol analyses were initially conducted just because subjects often use analogical problem solving and yet CASCADE did not model it, or at least, not very accurately (VanLehn and Jones, 1993b). However, it was discovered that different students used different strategies (or policies) for when to refer to the example and how much information to copy from it. The Good learners tend to use analogy less frequently and to copy less material than the Poor learners, who tend to copy the whole solution instead of trying to solve the problem themselves. This observation and other data suggest that maximizing one's use of analogy hurts learning and minimizing it help learning (VanLehn, 1995c).

At any rate, this review of the literature ended up breaking new ground in cognitive skill acquisition, and as a consequence took considerably longer than anticipated.

4 Modeling the major phenomena

We picked the expert-novice phenomena as our first group of phenomena for modeling, since we had access to the raw data for some of the seminal studies in this area (Chi et al., 1981; Larkin, 1983). Analysis of the expert protocols led us to the same hypothesis as Ericsson (albeit independently): experts plan absract solutions in their heads and novices do not. Koedinger and Anderson (1990) also reached the same conclusion for geometry expertise.

For instance, if a physics expert is given a problem that requires more than a dozen equations to answer algebraically, the expert will construct an abstract plan consisting of only one or two major equations (selected from a small set, including Newton's law, conservation of energy, conservation of momentum and kinematics). Constructing this plan seem to involve search, but not much search, as such plans are rather short. Once the plan is constructed, the expert implements it by writing the major equations and all the minor equations required to support the major ones.

We investigated several techniques for representing abstraction, starting with Koedinger's ideas about "diagram configuration schemas. Unfortunately, the abstractions that physics experts plan with are equations, and these can be applied in many ways, depending on what quantities are sought and given, what objects can be grouped together into a single object,

4

and what time intervals are relevant. None of the standard representations for planning operators had the flexibility that our task domain demanded.

Although one has the impression that an expert is simply applying a single planning operator each time they say something like "I'll use Newton's second law to get the acceleration," there is too much flexibility in the application of major principles to represent all the knowledge in a single structure. We found it necessary to use multiple rules that elaborate the problem description with possible objects, time intervals, forces, energy and momenta. The major principles are then applied to this elaborated description. We couldn't find a way to build the elaborative inferencing into the principles themselves, as is done in Koedinger's model and other planning formalisms.

Unexpectedly, this yields a much simpler account of the expert-novice shift. We had thought that experts could plan because they possessed abstract planning operators that novices lacked, and thus a major impact of learning was building such abstract operators. Because we ended up not using such operators, we have an even simpler account of learning. We assume that both experts and novices have the algebraic skill to apply algebraic principles abstractly, without writing them down or worrying about the mathematical details. We also assume that both experts and novices know the rules required for elaborating the problem descriptions (e.g., that a string exerts a tension force; that two abutting blocks can be considered a single objects, etc.). However, the experts' elaborative rules are so strong that they can apply them almost effortlessly, and conclusions produced by the rules are easily retained in memory. The novices, on the other hand, cannot do such elaborations in memory. They need to write elaborations down (e.g., as forces on a free-body diagram, or candidate equations). Moreover, they are likely to overlook possible elaborations. To put it a little differently, the experts' *see* much more in a problem than the novices. They see compound objects, forces, energies and other entities that are not mentioned in the problem statement. Novices have to work to "see" these entities, and even then, they may not see them all.

The implementation of this account progressed as far as building a problem solver that simulated an expert. Moreover, it could solve problems from all 6 textbook chapters on mechanics, whereas the original version of CASCADE was limited to 1 chapter. This work was reported in Jonathan Rubin's Masters Thesis (Rubin, 1994). Unfortunately, it was built on top of an older version of the CASCADE architecture, and would have to be modified in order to run on the latest version. We do not know if merely turning on the memory model and given CASCADE hundreds of hours of practice would cause an expert-novice shift, but it seems likely that it would.

5 Porting the learning events code

The last task of the project was to import the learning algorithms used in the original version of CASCADE. Cascade had three such algorithms: Explanation-based learning of correctness, analogical search control and analogical abduction (VanLehn and Jones, 1993a). Subsequent empirical work found no evidence for analogical search control (VanLehn and Jones, 1993b; VanLehn, 1995a). We ported the code for explanation-based learning, but have not yet ported the code for analogical abduction.

6 Final status

With only 15 months of funding, we were able to complete many of our goals but not all of them. We did derive a theory of cognitive skill acquisition that covers major phenomena in both the intermediate and late phases of skill acquisition. This theory is best expressed in a review article (VanLehn, 1995b), although because it is a review article, the theory is hidden between the lines. The basic assumptions of the theory are:

- All memory items have the same status; there is no procedural/declarative distinction.

- The only automatic, architecture mechanism are the traditional, basic ones from memory research: strengthening and retrieval cues that include context.

- All learning phenomena that are not due to strengthening are due to the learners' deliberate reformulation of their knowledge.

- Such reformulations are manifested as learning events, wherein learners interrupt their problem solving and reflect on the task domain knowledge itself.

This theory is considerably simpler than Act*, Soar and other theories. Because its model of memory is simpler, it is unable to model phenomena that only show up in brief tasks (e.g., semantic priming). However, its simplicity may allow it to model more phenomena at the molar scale that is most interesting to educators and others who are interesting in understanding learning in order to improve training.

Constructing this theory involved a considerable amount of protocol analysis. This was not anticipated in the original proposal. However, it led to three articles on the nature of learning events and analogical problem solving (VanLehn, 1995a; VanLehn, 1995d; VanLehn, 1995c).

The four articles just cited, along with a fifth one comparing our model of the self-explanation effect to 2 other models (VanLehn and Jones, 1995) and Rubin's Master's Thesis (Rubin, 1994), constitutes the main publications produced during this grant.

We developed parts of a computer model that embodies the theory. The parts include a cognitive architecture, a model of physics expertise and a version of the original CASCADE EBLC learning mechanism. Although it would take considerable work to complete the model, the graduate students and programmer who implemented the model have all left, so development on the model has halted temporarily. Nonetheless, the implementation taught us many things about cognitive skill acquisition, and these are reflected in the 6 publications mentioned above.

References

Anderson, J. R., Conrad, F. G., and Corbett, A. T. (1989). Skill acquisition and the LISP tutor. *Cognitive Science*, 14(4):467–505.

Anderson, J. R. and Fincham, J. M. (1994). Acquisition of procedural skills from examples. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 20(6):1322–1340.

Bovair, S., Kieras, D. E., and Polson, P. G. (1990). The acquisition and performance of text-editing skill: A cognitive complexity analysis. *Human-Computer Interaction*, 5:1–48.

Chase, W. and Simon, H. (1973). Perception in chess. *Cognitive Psychology*, 4:55–81.

Chi, M., Feltovich, P., and Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5:121–152.

Ericsson, K. A. and Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. This volume.

Koedinger, K. and Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14:511–550.

Larkin, J. (1983). The role of problem representation in physics. In Gentner, D. and Stevens, A., editors, *Mental Models*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95(4):492–527.

Priest, A. G. and Lindsay, R. O. (1992). New light on novice-expert differences in physics problem solving. *British Journal of Psychology*, 83:389–405.

Rubin, J. (1994). A model of expert problem solving in elementary mechanics. Unpublished Masters thesis.

Singley, M. and Anderson, J. (1989). *The transfer of cognitive skill*. Harvard University Press, Cambridge, MA.

VanLehn, K. (1989). Problem solving and cognitive skill acquisition. In Posner, M., editor, *Foundations of Cognitive Science*, pages 526–579. MIT Press, Cambridge, MA.

VanLehn, K. (1995a). Analogy events: How examples are used during problem solving. Submitted for publication.

VanLehn, K. (1995b). Cognitive skill acquisition. In Spence, J., Darly, J., and Foss, D. J., editors, *Annual Review of Psychology, Vol. 47*. Annual Reviews, Palo Alto, CA. in press.

VanLehn, K. (1995c). Looking in the book: The effects of example-exercise analogy on learning. Submitted for publication.

VanLehn, K. (1995d). Rule learning events in the acquisition of a complex skill. Submitted for publication.

VanLehn, K. and Jones, R. (1993a). Integration of analogical search control and explanation-based learning of correctness. In Minton, S., editor, *Machine Learning Methods for Planning*, pages 273–315. Morgan Kaufmann, Los Altos, CA.

VanLehn, K. and Jones, R. (1993b). Learning by explaining examples to oneself: A computational model. In Chipman, S. and Meyrowitz, A., editors, *Cognitive Models of Complex Learning*, pages 25–82. Kluwer Academic Publishers, Boston, MA.

VanLehn, K. and Jones, R. M. (1995). Is the self-explanation effect caused by learning rules, schemas or examples? Submitted for publication.

VanLehn, K., Jones, R. M., and Chi, M. T. H. (1992). A model of the self-explanation effect. *The Journal of the Learning Sciences*, 2(1):1–59.